

REMARKS

This paper is being provided in response to the October 20, 2004 Final Office Action for the above-referenced application. In this response, Applicant has amended Claims 1, 7, and 14 in order to clarify that which Applicant deems to be the claimed invention. Applicant respectfully submits that the amendments to the claims are all supported by the originally filed application.

With respect to the objection to the drawings, Applicants respectfully submit that the claims, as amended herein, are shown by the present drawings. Applicants note that the application as originally filed describes initiating a background copy for R1/R2 synchronization once communication between the R1 and R2 devices is resumed. See, for example, page 13, lines 10-14. Thus, the step 196 of the flow chart 50' of Figure 9 resumes communication between the R1 and R2 devices and, in accordance with the description on page 13, initiates synchronization of the volumes.

The rejection of claims 1-19 under 35 U.S.C. 102(b) as being anticipated by U.S. patent number 5,889,935 to Ofek, et al. (hereinafter "Ofek") is hereby traversed and reconsideration thereof is respectfully requested in view of amendments to the claims contained herein.

Independent Claim 1, as amended herein, recites a method of reversing a communication path between a first volume on a first storage device and a second volume on a second storage device. The method includes suspending communication between the first and second volumes while maintaining operations for other volumes of the storage devices, causing the first volume to change from a source volume to a destination volume without destroying the first volume, and

causing the second volume to change from a destination volume to a source volume without destroying the second volume. The method also includes, after causing the first volume to change from a source volume to a destination volume and causing the second volume to change from a destination volume to a source volume, initiating a synchronization of volumes to cause data to be copied from the source volume to the destination volume and prior to completion of the synchronization, resuming communication between the first and second volumes and resuming data access operations to the first and second volumes, where, in response to a data access operation to the second volume and valid data for the data access operation existing only on the first volume, the data access operation to the second volume is satisfied by accessing data from the first volume and where, in response to a write of particular data to the second volume, the particular data is transferred from the second volume to the first volume irrespective of whether the synchronization is complete. Claims 2-6 depend from claim 1.

Independent claim 7, as amended herein, recites a method of managing volumes on storage devices. The method includes receiving a command requesting reversal of a communication path between a first volume on a first storage device and a second volume on a second storage device, suspending communication between the first and second volumes while maintaining operations for other volumes of the storage devices, causing the first volume to change from a source volume to a destination volume without destroying the first volume, and causing the second volume to change from a destination volume to a source volume without destroying the second volume. The method also includes, after causing the first volume to change from a source volume to a destination volume and causing the second volume to change from a destination volume to a source volume, initiating a synchronization of volumes to cause

data to be copied from the source volume to the destination volume and, prior to completion of the synchronization, resuming communication between the first and second volumes and resuming data access operations to the first and second volumes, where, in response to a data access operation to the second volume and valid data for the data access operation existing only on the first volume, the data access operation to the second volume is satisfied by accessing data from the first volume and where, in response to a write of particular data to the second volume, the particular data is transferred from the second volume to the first volume irrespective of whether the synchronization is complete. Claims 8-13 depend from claim 7.

Independent claim 14, as amended herein, recites a computer program product that reverses a communication path between a first volume on a first storage device and a second volume on a second storage device. The computer program product is recited as including executable code that suspends communication between the first and second volumes while maintaining operations for other volumes of the storage devices, executable code that causes the first volume to change from a source volume to a destination volume without destroying the first volume, and executable code that causes the second volume to change from a destination volume to a source volume without destroying the second volume. The computer program product is also recited as including executable code that initiates a synchronization of volumes to cause data to be copied from the source volume to the destination volume after causing the first volume to change from a source volume to a destination volume and causing the second volume to change from a destination volume to a source volume and executable code that resumes communication between the first and second volumes and resumes data access operations to the first and second volumes prior to completion of the synchronization, where, in response to a data access operation to the second volume and valid data for the data access operation existing only on the first

volume, the data access operation to the second volume is satisfied by accessing data from the first volume and where, in response to a write of particular data to the second volume, the particular data is transferred from the second volume to the first volume irrespective of whether the synchronization is complete. Claims 15-19 depend from claim 14.

Ofek discloses two data storage systems that are interconnected by a data length for remote mirroring of data. Each volume of data is configured as local, primary in a remotely mirrored volume pair or secondary in a remotely mirrored volume pair. Figures 14 and 15 of Ofek and the corresponding text in Column 33 disclose data migration which may be needed when recovering from an all-link failure after continued processing upon a primary (R1) volume. Column 33 also discloses that data migration may also occur when a data center or host processor is moved from a local site to a remote site. Figure 15 of Ofek shows a first step where a host processing with the primary (R1) volume is suspended. Figure 15 shows a second step where the changed tracks of the primary (R1) volume, as indicated in a bit mask, are copied to the secondary (R2) volume. Following that step, Figure 15 shows a last step where the migration is finished so that the primary (R1) and secondary (R2) volumes are synchronized, and processing may resume using either of the volumes as a primary volume. In connection with this, Column 34 lines 54-63 disclose:

Once this copying is done, the migration task is finished. The primary (R1) and the secondary (R2) volumes are in sync, and they contain the same data. *Host processing may then resume* by accessing the primary (R1) volume in remotely mirroring data to the secondary volume. Alternatively, before resuming host processing, the linked data storage system could be reconfigured to reverse the roles of the primary (R1) and the secondary (R2) volumes, so that the host would directly access what was the secondary (R2) volume. (emphasis added)

Column 10, lines 39-47 of Ofek disclose:

Accordingly, each data storage device keeps data validity information about its mirrored device. If for some reason a device is not accessible, either the primary or the secondary device, every new write command goes to the accessible mirrored device along with information that the not accessible device has a track which is not valid. As soon as the non-accessible device becomes accessible, then automatically, as a background operation, the drives re-synchronize.

As indicated by the flowchart of Figure 15 and the text of Column 34, set forth above, Ofek discloses that the R1 and R2 volumes need to be synchronized prior to swapping the R1 and R2 volumes. That is, before the R1 and R2 volumes may be swapped in Ofek, the data on the R1 volume must be the same as the data on the R2 volume. This is because Ofek does not disclose a mechanism of resuming operation without first synchronizing the R1 and R2 volumes when the R1 and R2 volumes have been swapped.

Ofek also discloses in the flowchart of Figure 15 host processing is suspended (i.e., data accesses to R1 and R2 are suspended) prior to the synchronization and the swap and is only resumed after the synchronization and swap have taken place. Column 34, lines 50 and 51 state: “In step 480 of FIG. 15, the migration task *suspends host processing* with the primary (R1) volume.” (emphasis added). As set forth above, column 34, lines 54-59 state: “Once this copying is done, the migration task is finished. The primary (R1) and the secondary (R2) volumes are in sync, and they contain the same data. *Host processing may then resume* by accessing the primary (R1) volume in remotely mirroring data to the secondary volume.” (emphasis added). Thus, Ofek clearly contemplates a system whereby data accesses to R1 and R2 are suspended while the disclosed synchronization and swap are occurring. For the system

disclosed by Ofek, data accesses are not resumed until both the synchronization and swap have taken place.

Furthermore, the disclosure at column 10 of Ofek precedes the disclosure at column 14 and has nothing to do with swapping R1 and R2 as disclosed by Ofek at column 34. Instead, the disclosure at column 10 of Ofek teaches operations to be performed if one or the other of the primary or secondary (or links thereto) fail. The synchronization disclosed at column 10 of Ofek relates to the synchronization that takes place after the source of the failure is remedied.

Furthermore, it is respectfully submitted that there is no teaching in Ofek regarding allowing data accesses to R1 or R2 during the swap process or how to perform the swap process without first having to synchronize R1 and R2. The disclosure in column 10 of Ofek is for a relatively straight forward technique whereby the R1 accumulates information regarding data that has not been sent to R2 while the R1/R2 links are down. The disclosure in column 34 of Ofek relates to swapping R1 and R2, but appears to simplify the process by requiring 1) R1 and R2 be synchronized prior to the swap and 2) data accesses to R1 and R2 be suspended during the synchronization and swap. While this may make the swapping process of Ofek more simple, it suffers from the drawback of suspending data accesses by the hosts for the period of time while the synchronization and swap are occurring. This may be unacceptable in certain circumstances.

In contrast, the independent claims of the present application, as modified herein, specifically recite initiating a synchronization *after* swapping the first and second volumes (R1 and R2) and not *before* the swap, as taught by Ofek. In addition, the independent claims of the present application, as modified herein, specifically recite resuming communication between the

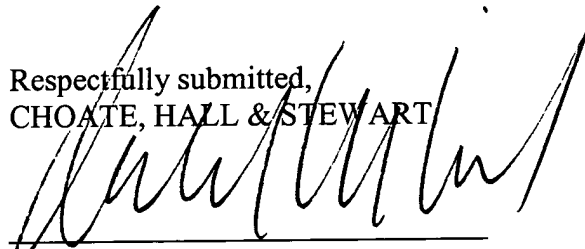
first and second volumes and resuming data access operations to the first and second volumes *before* completion of the synchronization, whereas Ofek specifically discloses suspending data access operations until *after* completion of synchronization. In addition, the independent claims of the present application specifically recite that, in response to a write of particular data to the second volume, the particular data is transferred from the second volume to the first volume *irrespective of whether the synchronization is complete* whereas Ofek specifically discloses *suspending communication* between R1 and R2 *until the synchronization/swap are complete* and discloses a specific mechanism for accumulating unsent data at R1 when there is no R1/R2 communication (i.e., setting bits of a table corresponding to unsent tracks as "invalid").

Thus, the present claimed invention provides the advantage of allowing communication between the R1 and the R2 devices to be resumed after the swap and before synchronization. In addition, data accesses (e.g., by hosts) are also resumed without first having to synchronize R1 and R2. This is shown in Figure 3 of the present application and described in the corresponding text beginning at the top of Page 15, which specifically discloses the situation where valid data may exist either on the local storage device or on the remote storage device. As set forth in Figure 3 and the corresponding description, if, after the swap has taken place, the valid data is located on the remote storage device and not the local storage device, the remote storage device is used to obtain the data. This feature of Applicant's present claimed invention addresses a drawback of Ofek where swapping entails suspending host access to R1 and R2 while synchronization and swapping are taking place. Applicants respectfully submit that this recited feature of the present claimed invention is neither shown, taught, nor suggested by Ofek.

Accordingly, for all of the reasons set forth above, Applicants respectfully request that this rejection be withdrawn.

Based on the above, Applicant respectfully requests that the Examiner reconsider and withdraw all outstanding rejections and objections. Favorable consideration and allowance are earnestly solicited. Should there be any questions after reviewing this paper, the Examiner is invited to contact the undersigned at 617-248-4038.

Respectfully submitted,
CHOATE, HALL & STEWART



Donald W. Muirhead
Registration No. 33,978

December 3, 2004
Date

Patent Group
CHOATE, HALL & STEWART
Exchange Place
53 State Street
Boston, MA 02109-2804
Tel: (617) 248-5000
Fax: (617) 248-4000